




INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification 6 : H04L 9/00		A1	(11) International Publication Number: WO 97/34391
			(43) International Publication Date: 18 September 1997 (18.09.97)
(21) International Application Number: PCT/US97/03816		(81) Designated States: AU, CA, JP, US, European patent (AT, BE, CH, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE).	
(22) International Filing Date: 12 March 1997 (12.03.97)			
(30) Priority Data: 08/615,534 12 March 1996 (12.03.96) US 08/679,863 15 July 1996 (15.07.96) US		Published With international search report. Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments.	
(71)(72) Applicant and Inventor: LEIGHTON, Frank, Thomson [US/US]; 15 Charlesden Park, Newtonville, MA 02160 (US).			
(74) Agents: JUDSON, David, H. et al.; Hughes & Luce, L.L.P., Suite 2800, 1717 Main Street, Dallas, TX 75201 (US).			

(54) Title: WATERMARKING PROCESS RESILIENT TO COLLUSION ATTACKS



(57) Abstract

The work to be protected by the watermark is first digitized (10), and then the baseline watermark vector is derived (12) and stored (14). Also, a watermark offset vector is created (16) and stored (18). The watermark offset vector and the baseline watermark vector are added together to generate a modified watermark vector (20). Next, the baseline watermark vector is replaced by the modified watermark vector in the digitized work to be protected (22). Finally, the watermarked work is returned to its original form (24).

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AM	Armenia	GB	United Kingdom	MW	Malawi
AT	Austria	GE	Georgia	MX	Mexico
AU	Australia	GN	Guinea	NB	Niger
BB	Barbados	GR	Greece	NL	Netherlands
BE	Belgium	HU	Hungary	NO	Norway
BF	Burkina Faso	IE	Ireland	NZ	New Zealand
BG	Bulgaria	IT	Italy	PL	Poland
BJ	Benin	JP	Japan	PT	Portugal
BR	Brazil	KE	Kenya	RO	Romania
BY	Belarus	KG	Kyrgyzstan	RU	Russian Federation
CA	Canada	KP	Democratic People's Republic of Korea	SD	Sudan
CF	Central African Republic	KR	Republic of Korea	SE	Sweden
CG	Congo	KZ	Kazakhstan	SG	Singapore
CH	Switzerland	LI	Liechtenstein	SI	Slovenia
CI	Côte d'Ivoire	LK	Sri Lanka	SK	Slovakia
CM	Cameroon	LR	Liberia	SN	Senegal
CN	China	LT	Lithuania	SZ	Swaziland
CS	Czechoslovakia	LU	Luxembourg	TD	Chad
CZ	Czech Republic	LV	Latvia	TG	Togo
DE	Germany	MC	Monaco	TJ	Tajikistan
DK	Denmark	MD	Republic of Moldova	TT	Trinidad and Tobago
EE	Estonia	MG	Madagascar	UA	Ukraine
ES	Spain	ML	Mali	UG	Uganda
FI	Finland	MN	Mongolia	US	United States of America
FR	France	MR	Mauritania	UZ	Uzbekistan
GA	Gabon			VN	Viet Nam

Watermarking Process Resilient to Collusion Attacks

TECHNICAL FIELD

The present invention relates generally to preventing unlawful copying of audio, video and other media that can be digitized and, 5 more particularly, to improved watermarking techniques that are robust even against multiple individuals who conspire together with independent copies.

BACKGROUND OF THE INVENTION

The proliferation of digitized media (audio, image and video) 10 and the ease with which digital files can be copied has created a need for copyright enforcement schemes. Conventional cryptographic systems permit only valid keyholders access to encrypted data, but once such data is decrypted there is no way to track its reproduction or retransmission. Such schemes thus provide 15 insufficient protection against unauthorized reproduction of information. It is known in the prior art to provide a so-called digital "watermark" on a document to address this problem. A "watermark" is a visible or preferably invisible identification code that is permanently embedded in the data and thus remains present within 20 the data after any decryption process. One example of a digital watermark would be a visible "seal" placed over an image to identify the copyright owner. However, the watermark might also contain additional information, including the identity of the purchaser of a particular copy of the material.

25 Many schemes have been proposed for watermarking digital data. In a known watermarking procedure, each copy of a document D is varied slightly so as to look the same to the user but also so as to include the identity of the purchaser. The watermark consists of the variations that are unique to each copy. The idea behind such

schemes is that the watermark should be hard to remove without destroying the document. Thus, a copy of a watermarked document should be traceable back to the specific version of the original from which it was created.

5 Although many prior art schemes claim to possess the "unremovable" property, all existing schemes are easily defeated by the following type of attack. Assume the attacker obtains two copies of the document that is being protected by the watermarking scheme. Each copy may have a different watermark, neither of
10 which is supposed to be removable. The attacker now makes a third version of the document (which he hopes will not have a traceable watermark) by averaging his two copies. For a pictorial document, for example, each pixel of the third version would be the average of the corresponding pixels in the watermarked copies.

15 Using existing approaches to watermarking, the third copy of the document produced by the attacker will look like the original versions but the watermark will be destroyed. This is because the "average" of two watermarks does not carry sufficient information to be tied to either of the watermarks individually. Thus, the
20 watermarking scheme can be rendered ineffective by simply averaging two copies of the document.

There is thus a need to devise a watermarking scheme that is immune to these and other such attacks, especially those in which the adversary obtains multiple copies of the original document.

25 **BRIEF SUMMARY OF THE INVENTION**

It is the principal object of the invention to describe a digital watermarking scheme wherein the watermark is robust against collusion by multiple individuals who each possess a watermarked copy of the data.

It is another object to describe such a scheme wherein the watermark cannot be removed by an adversary who obtains multiple copies of the original work.

It is a more general object of the invention to describe a 5 watermarking method that is secure against any form of attack including, without limitation, averaging attacks.

It is still a further object of the invention to describe a watermarking procedure wherein each of a set of copies of the work has a slightly-modified form of a "baseline" watermark that is placed 10 within a critical region of the data. The slight variations in the watermarks, however, are not perceptually visible and do not interfere with the works. If multiple persons collude to attempt to create an "illicit" copy of the work (i.e., a copy without a watermark), however, at least one of the modified watermarks is 15 present in the copy, thereby identifying both the illicit copy and the copier.

It is still thus another object to describe a watermarking scheme of the type recited above wherein combining copies of the same data set does not destroy the watermark.

It is a further object of the invention to describe such a 20 watermarking scheme that may be used to identify one or more of the parties who are colluding to destroy the watermark.

It is another more general object of the invention to describe a digital watermarking process that may be used as evidence in a Court 25 because it is robust against collusion.

According to the preferred embodiment of the invention, the work to be protected is digitized into a data file or string of data. A first digital watermark is then inserted in a first copy of the data file, preferably in a critical region of the data. A "critical" region may

consist of the entire document or alternatively will be some valuable portion of the work that will end up being significantly corrupted if the watermark is corrupted. A second digital watermark is then inserted in a second copy of the data file in a similar manner, and the 5 process is repeated for additional copies. According to the invention, the first and second digital watermarks are slight variations of a "baseline" watermark, which is kept secret, and one cannot perceive any differences between the first and second copies due to these variations. The baseline watermark may be a digital string that is 10 part of the original data being protected. Preferably, the variations are "randomized" in such a manner that if two persons were to collude to attempt to create an "illicit" copy of the work (i.e., a copy without a watermark), at least one of the first or second watermarks would still be present in the copy. After the watermark is inserted 15 into the work, the work can be converted back to its original form.

Thus, the scheme ensures that different possessors of watermarked copies of a work cannot create a "clean" copy that does not include at least one of the slightly-modified watermarks. Indeed, by comparing the watermark of the illicit copy with the 20 baseline watermark, one can determine the identity of the forger.

Although not meant to be limiting, preferably the "variations" are generated using a "random" offset, and in particular a "normal distribution."

BRIEF DESCRIPTION OF THE DRAWINGS

25 For a more complete understanding of the present invention and the advantages thereof, reference should be made to the following Detailed Description taken in connection with the accompanying drawings in which:

FIGURE 1 is a block diagram illustrating the method of inserting a digital watermark into a copy; and

FIGURE 2 is a block diagram illustrating the method for retrieving a digital watermark from a copy and correlating the retrieved

5 watermark with a stored watermark.

DETAILED DESCRIPTION

According to the invention, the work to be protected may be an image (photographs and graphics), video and/or audio (speech and music). The particular type of work is not relevant to the invention.

10 Referring now to FIGURE 1, the work, in whatever form, is digitized at step 10 into a data file or string of data either as part of the inventive technique or through some known A/D preprocessing. In the invention, there is a "baseline" watermark that is preferably stored and not used in making a particular copy of the work

15 (although this step is not necessarily required). This baseline watermark is then processed to create a set of one or more "modified" watermarks, each of which is related to the baseline watermark in a predetermined manner. Preferably, the "offsets" needed to create the modified watermarks are not fixed, however,

20 but are "randomized." In this way, a very small amount of "noise" is added to the offsets that does not alter the perception of the watermarked copies but still ensures that possessors of such copies cannot collude to remove all existence of the watermark in at least one illicit copy.

25 In general, collusion-type attacks are prevented according to the invention by constructing a watermark using randomness in a specific way. Preferably, an n-length digital string: x_1, x_2, \dots, x_n is derived at step 12 from the data to be watermarked and stored at step 14 for future reference. This may be referred to as the

"baseline" watermark. The string is preferably "critical" to the data in that corruption of the string will corrupt the data in a way that can be perceived and which will diminish the value of the corrupted document. Generation of the baseline watermark can be achieved in 5 many ways, e.g., by digitizing some portion of the document and using the resulting data or some subset thereof. (Whatever method is used is also used in the verification process, as discussed below). An n -length watermark vector w_1, w_2, \dots, w_n , is then created at step 10 16 and stored at step 18 for future reference. The vector is preferably created by choosing each w_i from a specified random distribution (preferably the normal distribution). The random distribution used for each w_i may or may not be the same (e.g., depending on whether it is desired to embed some specific serial 15 number data in the watermark). The watermark vector is then added at step 20 to the string x_1, x_2, \dots, x_n , and the result reinserted at step 22 into the original data to be protected. The work may then be converted back to its original form (image, video, audio, etc.) at step 24.

Assume it is now desired to retrieve the watermark from a 20 copy D' . This can be accomplished, as shown generally in FIGURE 2, by digitizing the copy D' at step 30 and then computing at step 32 the derived values x'_1, x'_2, \dots, x'_n , using the same algorithm used to compute the baseline watermark. Then, the method proceeds at step 34 by retrieving the original base line watermark, x_1, x_2, \dots, x_n , 25 from memory and subtracting out x_1, x_2, \dots, x_n from x'_1, x'_2, \dots, x'_n to compute a derived watermark w'_1, w'_2, \dots, w'_n at step 36. A correlation value (preferably an inner product) is then calculated between the derived watermark and w_1, w_2, \dots, w_n , retrieved at step 38, to produce a correlation value at step 40. The correlation value

is compared at step 42 to threshold levels, and if the correlation is high (step 44), then there is a match and a watermark is present. If the correlation is low (step 46), the watermark is not present. (The inner product scheme works by computing the absolute value of the sum $w_1w_1' + \dots + w_nw_n'$).

This scheme is immune to collusion because the watermark is random and because different watermarks are completely uncorrelated. In existing schemes, different watermarks are highly correlated and so it is easy for an attacker to exploit the correlation to destroy the watermark (e.g., by an averaging attack). In the invention method, there is simply not enough information contained in "t" different watermarked copies of the data in order for the adversary to remove the watermark. More specifically, if the attacker obtains "t" copies of watermarked data using the normal distribution to construct the watermarks (with watermarks $w_{11}, \dots w_{1n}$, through $w_{t1}, \dots w_{tn}$), it will appear to the attacker as if the original baseline watermark is $x_1 + (w_{11} + \dots + w_{1n})/t, \dots, x_n + (w_{t1} + \dots + w_{tn})/t$, which is not the true baseline watermark $x_1, \dots x_n$. The distinction is important since the former string is correlated with each of the watermarks $w_{11} \dots w_{1n}$ through $w_{t1} \dots w_{tn}$. In other words, the attacker simply does not have enough information in order to evade the watermark, no matter what sort of attack is used. Hence, one can prove that either the attacker must destroy the data or he must leave a trace of at least one of the component watermarks which will be revealed when the correlation test is run. Only someone with knowledge of the original baseline watermark could remove the watermark without detection.

Therefore, "m" copies of the work include variations of a baseline watermark such that up to "t" persons who possess those

copies cannot collude to create a "clean" copy (i.e., one without any watermark whatsoever). Stated another way, any "t" persons who collude in such a manner will always create an illicit copy that includes one of the modified watermarks. Comparison of the 5 watermark of the illicit copy with the baseline watermark then identifies what party made the copy (assuming there is a record of which party originally got which "version").

According to a preferred method, a first digital watermark is inserted in a first copy of a data file, preferably in a critical region of 10 the data. A second digital watermark is then inserted in a second copy of the data file in a similar manner, and the process is repeated for additional copies. As discussed above, the first and second digital watermarks are slight variations of a "baseline" watermark, which is kept secret, and one cannot perceive any differences 15 between the first and second copies due to these variations. Preferably, the variations are "randomized" in such a manner that if two persons were to collude to attempt to create an "illicit" copy of the work (i.e., a copy without a watermark), at least one of the first or second watermarks would still be present in the copy. In the 20 preferred embodiment, a watermark consists of a sequence of real numbers $W = w_1, \dots, w_n$, where each value w_i is chosen independently according to $N(0,1)$ (where $N(\mu, \sigma^2)$ denotes a normal distribution with mean μ and variance σ^2). The watermark may consist of a number (e.g., 1000) of randomly generated numbers 25 with a normal distribution having zero mean and unity variance. Alternatively, w_i could be selected according to $N(\mu_i, \sigma_i)$ where μ_1, \dots, μ_n can be a serial number corresponding to the copy being watermarked (or other information that may be embedded).

In order to detect the presence of a watermark W in a derived watermark signal W' , we preferably use a correlation function $\text{cor}(W, W') = |W \cdot W'|$, which is the inner product of two vectors. If W were selected according to the normal $N(0,1)$ distribution and W' is uncorrelated to W (but of the same order), then the correlation will be small (about \sqrt{n}). If W' is closely correlated to W , then the correlation will be large (about n). If W' is uncorrelated to W but is of a larger order (e.g., due to intentional or unintentional noise or attempts to hide the watermark), then the correlation might also be large. (Specifically, if W' is uncorrelated to W but has B times the magnitude, then the correlation is about $B \sqrt{n}$. If B is large, then the data D' will not resemble D . (The notion of large in this context depends on the application and the level of security/clarity desired). In any event, the watermark is said to be present if $\text{cor}(W, W') > c \sqrt{n}$, where c is a predetermined constant that depends on the application and level of security desired (e.g., $c = 4$).

The correlation will be low if the watermark is not present and the work is not destroyed. The correlation will be high if D' was derived from the watermarked document or if the data has been corrupted beyond recognition (the latter condition being determined by inspection).

As noted above, it is preferable that each of the "modified" watermarks be placed in a critical region of the data. Of course, the exact location will depend on the nature of the work being protected. It is also helpful if every entry in this region of data is largely uncorrelated with the other data. It has been suggested (by Cox et al) that this can be accomplished by embedding a watermark in the spectrum of an image, the temporal frequency domain of an audio

signal, or the spatio-temporal frequency domain of a video sequence. Although the above techniques are preferred, one may even encode the watermark in other less, desirable places (such as in the low order or least significant bits) of the data and still obtain the 5 advantages of the collusion-resistant feature of the invention where multiple parties may collude to remove the watermark.

Variations

In the embodiment discussed above, the original document (or an original baseline watermark vector) is stored in order to determine 10 whether the watermark is present in a copy of the document. In the embodiment previously described, the original baseline watermark vector is retrieved at step 34 and subtracted from the derived baseline watermark vector to produce the derived offset watermark vector. This step can be omitted without changing the detection 15 protocol or its results. In particular, the derived offset watermark vector may be set equal to the derived baseline watermark vector. This change increases the noise level in the correlation test, but not beyond tolerable levels. Further, the noise levels can be reduced by specially selecting the original offset watermark vectors to have low 20 noise (e.g., by selecting them to be orthogonal to the original baseline watermark vector to which they are being applied) or by running the correlation test on only specific components of the vectors.

Another improvement would be to remove the need to store 25 the original offset watermark vector. As discussed above, in one embodiment of the invention it is necessary to store a copy of the original offset watermark vectors (see, e.g., step 18) so that they can be later retrieved and correlated with the derived offset

watermark vectors (see, e.g., step 38). This step can be largely omitted by the following process.

The original offset watermark vectors are computed using a secret random hash function H . The function H maps copyright and other information that the user desires to embed in the document (e.g., "This picture is the property of XYZ Corp., unauthorized copying is forbidden") to the sequence of numbers $W = w_1, \dots, w_n$ that was used as the original offset watermark vector. The sequence of numbers preferably has same structure and function as discussed above and appear to be random, but the sequence is easily reconstructed given the secret function H and the underlying information to be inserted into the document. Hence, a watermark is identified by reconstructing the original offset watermark vector locally instead of retrieving the vector from a database. More generally, the text to be embedded may be a simple serial number, and this serial number can be retrieved from the document by checking all possibilities to see if there is a correlation. This check can be done locally if H is available, since all relevant original offset vectors can be regenerated as needed.

Thus, according to this variation of the present invention, one need not subtract the original picture before carrying out the dot product form of the correlation test described above in the main embodiment. In such case, the correlation test generates the old dot product (which is large, precisely what is desired) plus the dot product of the offset vector and the original picture. Since the offset vector is random, this dot product is small (in the noise range) for any picture. Therefore, one does not need the original picture to do the correlation test. Moreover, by using the secret random hash function H , one need not store the offset vectors. The function

maps a copyright notice or text into a sequence of independent Gaussian offsets (i.e., an offset vector). Then, one may choose the offset vector for some text to be $H(\text{text})$. Now, one need only remember the text, not the whole offset vector. The text may be 5 timestamped so that the same offset vector is only used once, although one can use the same offset vector more than once.

This method is provably secure, even against colluders, but has low memory requirements. A two-tiered version, wherein there are two hash functions (one for the sign and one for the magnitude 10 of the offsets) might be used as well. In this way, one of the two (sign or magnitude) would be kept in reserve and not released, even in the secure software.

The above-described variants can be combined advantageously to provide a scheme to prevent unauthorized copying of certain 15 media such as CD's and VCR videos. In this application, a given text -- such as "Do not copy" -- is used as the watermark. A VCR can then check for the presence of this watermark before allowing the copying to take place. This would be achieved by having the secret function H embedded in the VCR software or hardware in a secure 20 fashion, e.g., through a secure chip or via a protected software encryption scheme. The value of H would also be embedded securely in the hardware or software that generates the watermarked copy in the first place instance.

In the VCR/CD application, it may only be necessary to use a 25 single watermark for many copies of the document, in which case it may only be necessary to use a single watermark offset vector (e.g., H ("Do not copy")) for different documents. In this variant, the system must be secure against a different kind of collusion; namely, one in which the same watermark is used with different documents

instead of the case where the same document is used with different watermarks. Fortunately, the same analysis applies to both scenarios equally well, such that either scheme is secure against collusion.

In the above-described variant, the hardware/software that creates the watermarks is in secure hands (so that H remains secret and cannot be misused). For example, if the adversary is allowed to watermark a blank document, then the scheme can lose security. There are several ways, however, that security can be enhanced as is now explained.

10 In one approach, it is assumed that each copy of the watermarking software produces watermarks unique to the copy. For example, the XYZ Corporation watermarking software produces watermarks of the form $H(XYZCOPR | \text{Do not copy})$. Then, only the watermarks produced by that software would be compromised if the 15 XYZ software were stolen. (For simplicity, each version of the software could be the same except for a special key unique to the version.) Alternatively, the original offset watermark vectors can be derived as a function of the document that is being watermarked in addition to the text that is being embedded into the document. This 20 has the effect of making watermarks corresponding to "Do not copy" be different for each document in which they appear. For example, one might use $H(x_1...x_n | \text{Do not copy})$ as the original offset watermark vector for a document with features x_1, \dots, x_n into which the "Do not copy" text is embedded. Even further, the string 25 x_1, \dots, x_n may include random numbers so that offset vectors can be further differentiated in an effort to prevent attacks.

In order to confirm the presence of a watermark in the preceding examples, one still needs to know (or guess, perhaps by exhaustive search) the underlying text that was used to generate the

original offset vector. This process can be simplified by embedding serial numbers instead of text. Once the serial number is retrieved, a global database is consulted to find out what the text is. However, it is still necessary to be careful how a serial number is embedded since

5 exhaustive search over a space of 12-digit numbers would be costly and difficult. In such a case, it would be much better to separately embed say four (4) serial numbers, each with 3 digits. (Of course, such numbers and their characteristics are merely exemplary). Then, one would only have to search over a space of 1000 numbers
10 (instead of 1,000,000,000,000 numbers) four times. (This technique makes use of the fact that the watermarking procedures can be used to embed more than one watermark in a document.)

It is also possible to make the watermarking process more resilient to noise as well as more secure. This is achieved as follows.

15 Suppose that one desires to embed the text "Do not copy" in a document. Another good way of doing this is to embed multiple offset watermark vectors in the document. For example, we could use $H(y_1 | \text{Do not copy})$, $H(y_2 | \text{Do not copy})$, ..., $H(y_m | \text{Do not copy})$ for different values of y_1 , ..., y_m as the vectors. If any of
20 the watermarks is detected, then copying would not proceed. Such a scheme is more robust since all m vectors would have to be ruined by noise or be removed by an adversary before copying could proceed. If there is a chance p of being able to remove any one of the vectors, then the chance of losing all m is p^m (assuming independence), which is very small (e.g., if $p = .01$ and $m = 4$, then
25 $p^m = 10^{-8}$).

IN THE CLAIMS

1. A watermarking method, comprising the steps of:
 - (a) generating a digital string from the work to form a baseline watermark;
 - 5 (b) generating a set of watermarks each having a predetermined relationship to the baseline watermark; and
 - (c) inserting a respective one of the set of watermarks into a respective copy of the work to create a watermarked copy uniquely identified by the respective one of the set of watermarks; and
- 10 (d) repeating step (c) at least m times to create a set of m watermarked copies, each having a different one of the set of watermarks therein, such that if a subset of said m watermarked copies are averaged to produce an illicit copy of the work, at least one of the set of watermarks is detectable in the illicit copy.

15

2. The method as described in Claim 1 wherein the predetermined relationship is a set of random offsets.
3. The method as described in Claim 2 wherein the random offsets have a normal distribution having zero mean and unity variance.
- 20 4. The method as described in Claim 1 wherein the work includes an image.
- 25 5. The method as described in Claim 1 wherein the work includes an audio signal.

6. The method as described in Claim 1 wherein the work includes a video signal.

7. The method as described in Claim 1 wherein each of the 5 watermarks is inserted in a critical region of the digital data file.

8. The method as described in Claim 1 further including the step of comparing the watermark in the illicit copy with the baseline watermark to determine which possessor of a copy of the work 10 created the illicit copy.

9. A method of securing a work against copying, comprising the steps of:

(a) generating a set of watermarks each comprising a 15 vector of randomly-generated numbers; and

(b) inserting a respective one of the set of watermarks into a respective copy of the work to create a watermarked copy uniquely identified by the respective one of the set of watermarks; and

(c) repeating step (b) at least m times to create a set of m 20 watermarked copies, each having a different one of the set of watermarks therein, such that if a subset of said m watermarked copies are averaged to produce an illicit copy of the work, at least one of the set of watermarks is detectable in the illicit copy.

25 10. The method as described in Claim 9 wherein the work includes an image.

11. The method as described in Claim 9 wherein the work includes an audio signal.

12. The method as described in Claim 9 wherein the work includes a video signal.

5 13. The method as described in Claim 9 further including the step of comparing the watermark in the illicit copy with the set of watermarks to determine which possessor of a copy of the work created the illicit copy.

10 14. A method of protecting a work against illicit copying, comprising the steps of:

(a) generating a set of watermarks each having a predetermined relationship to a first watermark for the work; and

15 (b) inserting a respective one of the set of watermarks into a respective copy of the work to create a watermarked copy uniquely identified by the respective one of the set of watermarks; and

(c) repeating step (b) at least m times to create a set of m watermarked copies, each having a different one of the set of watermarks therein, wherein averaging a pair of said m watermarked

20 copies generates a copy of the work in which at least one of the set of watermarks can be detected.



15. The method as described in Claim 14 wherein the baseline watermark is derived from the work.

16. A method of generating secure copies of a document, comprising the steps of:

- (a) generating a set of watermarks each comprising a vector of randomly-generated numbers with a normal distribution having zero mean and unity variance; [and]
- 5 (b) inserting a respective one of the set of watermarks into a respective copy of the document to create a watermarked copy uniquely identified by the respective one of the set of watermarks; and
- 10 (c) repeating step (b) at least m times to create a set of m watermarked copies of the document that are secure against illicit copying.

17. A method of generating secure copies of a document, comprising the steps of:

- (a) generating a set of watermarks each comprising a vector of randomly-generated numbers;
- (b) inserting a respective one of the set of watermarks into a respective copy of the document to create a watermarked copy uniquely identified by the respective one of the set of vectors; and
- 20 (c) repeating step (b) at least m times to create a set of m watermarked copies of the document that are secure against illicit copying.



INTERNATIONAL SEARCH REPORT

International application No.
PCT/US97/03816

A. CLASSIFICATION OF SUBJECT MATTER

IPC(6) :H04L 9/00

US CL : 380/54

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

U.S. : 380/54, 4, 5, 23, 25, 49, 50

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 5,488,664 A (SHAMIR) 30 January 1996, see Abstract.	1-17
A,P	US 5,530,759 A (BRAUDAWAY ET AL) 25 June 1996, see Abstract.	1-17

<input type="checkbox"/>	Further documents are listed in the continuation of Box C.	<input type="checkbox"/>	See patent family annex.
--------------------------	--	--------------------------	--------------------------

•	Special categories of cited documents:		
"A"	document defining the general state of the art which is not considered to be of particular relevance	"T"	later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
"E"	earlier document published on or after the international filing date	"X"	document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
"L"	document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)	"Y"	document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
"O"	document referring to an oral disclosure, use, exhibition or other means	"Z"	document member of the same patent family
"P"	document published prior to the international filing date but later than the priority date claimed		

Date of the actual completion of the international search	Date of mailing of the international search report
09 JULY 1997	04 AUG 1997

Name and mailing address of the ISA/US Commissioner of Patents and Trademarks Box PCT Washington, D.C. 20231 Facsimile No. (703) 305-3230	Authorized officer <i>Deanie Goodenix</i> BERNARR EARL GREGORY Telephone No. (703) 306-4153
---	---